

## FUNCTIONAL ENTERPRISE BEAN

## TECHNICAL FIELD

5       The present invention is related in general to component-based multi-tier computing methods, and in particular, to a functional enterprise bean component, and a system for efficient management of components.

## BACKGROUND

10      With the rise of interconnected computer networks such as intranets and the Internet, complex transaction-based applications that are distributed over several networked computers became a possibility. In general, these transaction-based applications function in the following way. A software application program, which executes on a computer called a client, initiates a transaction that requires access to services provided by a distant computer, called a server. Examples of these services could be an update to a database such as a bank's database, execution of a purchase order such as in the case of purchase of a security, and the like. Typically, the client sends a "request" message to the server, which then sends a "response" message containing a response to the request.

15      Assume that the server is not a single computer, and instead, a collection of interconnected heterogeneous computers. The request message must then be formatted in such a way that all the interconnected computers can understand and respond to the request message. This is generally the case where a client coordinates a distributed transaction, as is the case in a typical client-server model. If the collection of interconnected computers is configured in an object-oriented programming model, then software objects (or objects) that 20     are capable of working together to provide a response to the request message can be distributed among the several computers. But in order to access the objects from a remote computer such as a client, the objects must somehow publish their existence, their addresses, their properties, the services they provide, and other details to the "outside" world. Then, a 25     client may be able to use the services provided by sending a request message in a manner similar to making a remote procedure call ("rpc") and obtaining a response to that message.

30      Three paradigms arose as a result of the need to standardize the methods by which

objects could be distributed and accessed over a network. These are Microsoft Corporation's Distributed Component Object Model (DCOM), JavaSoft's Java/Remote Method Invocation (Java/RMI), and Object Management Group's Common Object Request Broker Architecture (CORBA).

5        Though differences exist among these models, they principally work in the following way. Objects that provide services are typically located on servers. These objects are queried by applications running on clients using a specified data communication transport layer protocol—the Object Remote Procedure Call (ORPC) for DCOM; the Java Remote Method Protocol (JRMP) for Java/RMI; and the Internet Inter-ORB Protocol (IIOP) for  
10      CORBA. A client suitably formats a query message in the appropriate protocol language and transmits the query message, which is routed to the appropriate server, whereupon it is examined, and a response message is formatted and routed back to the client.

15        As used in the present application, the term "object" may mean the object definition, associated operations, attributes, as well as the implementation for that object. Persons of ordinary skill in the art understand that the term "object type" is used to refer to the definition of the operations and attributes that software external to the object may use to examine and operate upon the object. An object of a given type can support multiple interfaces. Additionally, the term "object" may sometimes be used to refer to an actual run-time instance  
20      of an object and will be made clear by the context.

25        As stated above, clients should be configured to understand the services offered by the various objects located at the servers. In DCOM, this is achieved by exposing certain methods of a DCOM server. In this paradigm, a client acquires an interface pointer to one of the server's exposed methods. This interface pointer presents itself to be locally addressable to the client. An object at the client—for example, a proxy object—uses the interface pointer and calls the server's methods in order to determine their structure and properties.

30        In Java/RMI, objects are encoded and transmitted ("marshaled") in streams of bytes by a process called "serialization." A server configured to be a Java/RMI server comprises objects that have predefined interfaces, which can be used to access the server objects remotely from another machine's Java Virtual Machine (JVM). A Java/RMI server object interfaces declare a set of methods that indicate the services offered by that server object. A program resident on the server called an RMI Registry stores and makes available to clients

information about available server objects. Typically, a client object obtains information regarding the methods and other properties of a server object by performing an operation such as "lookup" for a server object reference. This lookup typically works by the client object specifying an address in the form of a Universal Resource Locator (URL) and transmitting the address to the server's RMI Registry.

During the early 1990s, there had been a shift from a two-tier, client-server application model to a more flexible architecture that incorporated modular separation of concerns. This resulted in breaking the application model to include a three or more-tier models, which separated business logic from system services and user interfaces. The business logic is placed in a different layer-frequently called the middle-tier or middleware-and included transaction monitoring, messaging, object request brokering, and other services. This also allowed the deployment of lightweight clients that are easy to deploy.

In general, these architectures are thought to simplify developing, deploying, and maintaining enterprise applications. The cited advantages for this architecture included a focus on the specifics of business logic programming without the need for transaction management and user-interface considerations. Thus, the multi-tier architecture enabled developers to focus on the specifics of programming their business logic, relying on various back-end services to provide the infrastructure, and client-side applications (both standalone and within web browsers) to provide the user interaction. Because of the logical separation achieved by the multi-tier model, business logic can be developed once and deployed on servers appropriate to the needs of an organization in such a way that the servers can be readily scaled to meet changing business conditions.

While the above methods are directed toward distributing the objects over a collection of networked computers, there still was a need to develop a standardized way to develop transaction-oriented systems that are vendor-independent. Traditionally, products such as TUXEDOT<sup>TM</sup>, CICS<sup>TM</sup>, have been used to provide transaction management, security, concurrence, persistence and other services in transaction-oriented applications such as travel planning, banking and stock market transactions. But there was a problem of integrating these products with the newly emerging distributed object computing models discussed above. Additionally, it was determined that transaction-related programming logic should advantageously be separated from business logic. The middle tier, i.e., the layer between the

client and the “back-end,” is therefore modeled so that transaction monitors are separated from programming entities that perform only the business logic. It should be noted that the transaction monitors are infrastructure components insofar as they are a part of application server architecture. Transaction monitors need not be implemented as objects.

5        This need for a standardization resulted in the Enterprise JavaBean™ (EJB) computing model, which took advantage of the technological advances such as the JDBC technology. The servlet technology enabled developers to create CGI-like behaviors that could run on any web server that supported the Java platform. The JDBC™ technology provided a model for integrating the “Write Once, Run Anywhere™” features of the Java  
10      programming language to existing database management systems. And the JavaBeans component architecture enabled encapsulation of complete sets of behavior into easily configurable, readily reusable components on the client side. The convergence of these three concepts-server-side behaviors written in the Java programming language, connectors to enable access to existing enterprise systems, and modular, easy to deploy components-led to  
15      the development of the EJB standard. It should be noted that the Enterprise JavaBeans might be used in conjunction with any of the technologies described herein.

In the discussion that follows, a reference to a “bean” indicates it is an Enterprise JavaBean or a similar component. Additionally, an “application server,” a “middleware server,” an “enterprise application server,” or a “server” are terms used interchangeably unless otherwise qualified, each of which terms indicates a computer connected to a communication network and at least one of a plurality of back-end computers or data repositories, for example, a legacy database server.  
20

The EJB is a server-side component model that incorporates the business logic in software components called beans. An enterprise application server or simply, a “server”  
25      that is configured as an EJB server is responsible for managing services such as transactions, persistence, concurrence and security.

The EJB architecture, which is a distributed object architecture, comprises one or more computers configured as EJB server(s), EJB container(s) that run within the EJB server(s), beans that execute within these containers, and EJB clients. EJB clients locate  
30      beans using an interface called the Java Naming and Directory Interface (JNDI); and utilize transaction support services such as Java Transaction Service (JTS). Typically, the beans are

similar to the server objects described above, except that they contain only business and no transaction-monitoring logic. EJB servers provide certain predefined system services such as raw execution environment, multiprocessing, load balancing; access to devices, naming and translation service provisioning, and making EJB containers visible. EJB containers provide the interface between the EJBs and EJB clients. In general, an EJB client can access a bean only via EJB container-provided methods, which in turn invoke the bean's methods.

An EJB object is a "remote" interface object that implements the remote interface of a bean. It wraps a bean instance on the server and expands its functionality to include transactions, security and other system-level operations to the bean at runtime. Typically, an EJB container vendor provides the utilities to create the EJB object.

Similar to the EJB object class, the EJB home class is also automatically created when an Enterprise JavaBean is installed in a container. The EJB home class implements all the methods defined by the bean's home interface and is responsible for helping the container manage the bean's life cycle. The EJB home is responsible for locating, creating and removing enterprise beans, by working with the EJB Server's resource managers, instance pooling, and persistence mechanisms as needed transparently to the developer.

In the EJB model, there are two types of beans-session beans and entity beans. Session beans, as their name indicates, are used to keep track of and manage a client's session with the server. A session bean is usually associated with a single EJB client. In general, session beans do not survive a server shutdown or crash. Session beans can preserve their states through a user's session lifetime.

A stateful session bean contains state information. Typically, when a client initiates a connection with the server, the client specifies the type of EJB to which it connects on the server. When the client is finished with the session, the stateful session bean is typically removed. Thus, in a typical scenario, only one client can use an instance of a particular stateful session bean at a given time.

A stateless session bean instance is an object that can be reused by different clients, one at a time. Since it can be reused, a stateless session bean instance need not be destroyed until all clients are finished with them. Additionally, it contains no state information, and therefore, it is impossible for the client to tell one stateless session bean from another of the same type. Thus, one stateless session bean of a particular type is virtually indistinguishable

from another of the same type. In view of these properties, stateless session beans of each type are pre-created by an EJB container in sufficient numbers and placed together in pools of stateless beans. Whenever a client requests a stateless bean of a certain type, a bean of that type is removed from the proper pool and given to the client. Once the client is done, the  
5 bean is returned to the pool. These beans are only created and/or deleted by the EJB container itself.

Multiple EJB clients, on the other hand, can share entity beans. They always have states and these states persist server shutdown or crash. Typically, these are implemented by writing any state-related data into a persistent store so that the bean can be reconstructed  
10 from the stored data after a server is restarted subsequent to a shutdown or a crash.

An entity bean instance is usually associated with a specific row (record) in a specific table. More complex mappings are also possible. An entity bean is also associated with a transaction. The bean's state is always in step with the transaction—if the transaction commits, the bean's state persists, and if the transaction rolls back, the bean is brought back to its initial state. This synchronization is automatic and seamless to a developer, unless the bean is deployed as having bean-managed persistence, or no transactional persistence at all.  
15 In general, a client can access a particular entity bean instance using an identifier or a key.

#### SUMMARY

As stated above, entity beans usually provide a mapping to data elements such as a  
20 row in a database, and therefore they are widely used to model business data. Since session beans provide session services, they are used to model business functions. Where transactional synchronization is needed, for example, in the case of a trouble-ticket entry system, the traditional method is to use stateful session beans together with SessionSynchronization interface as method of implementing such synchronization. But this  
25 method of implementation involves the generation of a large amount of code, and therefore is not elegant. Currently the number of clients, and not the number of resources limits the number of session beans.

Further, the traditional approach recommends that an entity bean be accessed only by a session bean and not directly by a client application according to the following model:

30 Client --> Session bean --> Entity bean --> Database Record

This is based on the idea that the entity bean must model only entities (i.e., one or more database records). A problem with this approach is that the service providers are not visible to the client, where there could be a need for such visibility. As stated earlier, if there is a need for a number of database records to be simultaneously accessed by a client, then  
5 there will be a proliferation of such beans if each record is modeled as an entity bean. In extreme cases, such a proliferation may result in a depletion of available resources for other tasks. Further, since access to these entity beans should be serialized, each successive client with a need to access a particular row (or an entity bean) for an update should wait for its turn. Moreover, there is an implicit assumption that the client should know beforehand the  
10 primary key (or the identifier) of the entity bean it needs to contact. The current methods do not provide an easy solution to these problems, though one could craft an unwieldy solution using session beans. Accordingly, there is a need for a new bean type that provides transactional capability of an entity bean without its persistence, and models business functions. These capabilities are transparent to the clients that request the bean's services.

15 Transactional synchronization as well as transactional integrity can be achieved by implementing a system using entity beans rather than session beans. Noting that object oriented computing methods departed from a procedural or a functional programming to encapsulate data rather than functions in objects, it has been discovered that mapping an object to a function, rather than to a data element provides significant advantages, especially  
20 in managing limited resources that are accessed by a plurality of clients. Further, the presently available Enterprise JavaBeans (EJB) can be programmed to achieve such a mapping.

Accordingly, in an embodiment, the disclosure is directed to a novel "functional" bean, which is devoted to modeling a business function. Clients do not need to know the  
25 particular primary key or identifier as in the case of an entity EJB; rather a client knows only a well-known Service Manager bean to obtain a handle to the correct type of functional bean.

If a client needs to request a service offered by such a functional bean, it can be invoked directly. In one aspect, these functional beans are created and managed by a bean container such as an EJB container. The container controls access to these beans by way of a  
30 Service Manager. In one embodiment, the Service Manager itself is modeled as a functional bean, whose function is that of a manager of the other functional beans. Thus, in one

embodiment, a three-tiered model that can accomplish the task contains the following structure.

Client --> Functional bean --> Database view/record

There could be cases where a system does not require some attributes such as state management and persistence. In such a situation, a system designed using functional beans can be easily implemented without these attributes, typically associated with an entity EJB.

In one aspect, what is disclosed is a computer system comprising a plurality of sets of functional beans each set comprising at least one functional bean assigned to perform a particular business method, the computer system comprising: a microprocessor; a memory

device coupled to the microprocessor; a service manager program coupled to the memory device and configured to a number of receive requests for at least one of a plurality of types of transactions from a plurality of clients; a load-sharing program coupled to the service manager program and configured to create instances of functional beans based on the number of requests from the plurality of clients; the service manager program configured to obtain a handle to an instance of a functional bean based on a type of transaction requested by a client; and the service manager program configured to return the handle to the client, wherein the client uses the handle to interact with the functional bean to execute a business method.

In another embodiment, the present disclosure includes a computer-processor-executable software code stored in a computer-readable memory, said code comprising:

instructions to direct a computer processor to receive a first request from a client for a handle to a particular type of functional bean, said functional bean which comprises code to execute a particular business function; instructions to direct the computer processor to create an instance of a functional bean of the particular type requested; instructions to direct the computer processor to obtain a handle to said instance of said functional bean; and instructions to direct the computer processor to transmit to the client said handle to said instance of said functional bean, wherein the computer processor, responsive to a second request from the client enables the client to execute code comprised in the functional bean to accomplish the particular business function. In further aspects, the code further comprises instructions to receive the first request and the second request from the client via a computer network; or instructions to create a number of instances of functional beans of the particular type, said number being dependent on availability of resources; or instructions that allow a

functional bean to instantiate a second bean of a second type in order to execute the business logic contained in the second bean instance. In a yet another aspect, the code further comprises instructions that allow an instance of a session Enterprise JavaBean to invoke the business methods contained in the functional bean. In a yet further aspect, the functional  
5 bean can call any type of EJB beans, including entity beans.

#### BRIEF DESCRIPTION OF THE DRAWINGS

These and other objects, features and advantages of the present invention can be more readily understood from the following detailed description with reference to the  
10 accompanying drawings, where like numerals represent like parts, and in which:

Figs. 1(a)-(f) depict several exemplary architectural views of a client transaction utilizing a functional bean;

Fig. 2 shows a method of implementing the invented functional bean;

Fig. 3 depicts an exemplary method by which a functional bean obtains a handle to an  
15 external resource;

Figs. 4(a)-(b), depict views of hardware architecture comprising a plurality clients coupled to a network, to which is communicatively coupled an enterprise application server;

Fig. 5 depicts an architecture of software executing on an exemplary enterprise application server;

Fig. 6 shows an EJBObject acting as a proxy to mediate access to an instance of a functional bean;

Fig. 7 is a flow-chart depicting the steps included in an illustrative transaction using a functional bean of the present invention; and

Fig. 8 is an alternative embodiment including a web-based client program executing a  
25 transaction.

## DETAILED DESCRIPTION

### “Functional” EJB

An aspect of the present invention includes incorporating certain functional (i.e.,  
5 functions as objects) programming methods in the Enterprise JavaBean (EJB) model to  
accomplish certain tasks that are not easily accomplished by the session or entity beans as  
they are currently defined. It has been discovered that certain tasks such as updating a  
customer account are more advantageously implemented using a functional approach, which  
enables a programmer to define the business logic in the form of a series of steps performed  
10 on a set of entities (or objects).

This functional approach is advantageously accomplished by creating a novel server-  
side EJB. This novel EJB is hereinafter referred to as a “functional EJB” or a “functional  
bean.” It should be noted that the principles of the present invention could also be used to  
enhance other server-side or client-side objects such as JavaBeans, or other entities, thereby  
allowing a programmer to take advantage of the functional bean type in a true object-oriented  
15 manner.

In an embodiment, the characteristics of the functional EJB designed according to the  
present invention include:

- a) Transactional awareness, which is a default characteristic in an embodiment and is  
20 implemented by providing a persistent database to which transactional data are written.
- b) Optionally, transparent support of transactional persistence, i.e., the data objects  
operated on by a functional bean’s methods and any state information are persisted.
- c) Allowing multiple clients to use a single instance of a functional bean. Since  
functional beans are designed to incorporate functions rather than data, they can receive a  
25 client request from a queue and service the client request without changing the request’s own  
internal state. These requests can be serviced according to a particular scheme, such as a  
first-come, first-serve, and the like. In alternative embodiments, the functional bean can  
service client requests using a priority method as designed by a programmer or others skilled  
in the art.
- 30 d) A pool of functional beans of each functional type is maintained by the enterprise  
application server to load-balance the requests from clients.

e) Load sharing-handled according to a strategy that is either static or dynamic, such as round robin or least busy. A client need not know which instance of each type of functional bean is in communication with the client at any given moment; the client is required only to be aware of the type of the bean with which it is communicating.

5 f) The number of functional beans of each type in the pool can be configured at startup and/or dynamically controlled by the bean's container. Depending on the number and availability of limited resources—for example, the number of open connections allowed by a legacy system—determines the number of bean instances of each type. In an embodiment, a system administrator makes an appropriate choice of the number of instances 10 of each bean type either statically or dynamically.

h) The functional bean does not map to one or more rows in a database.

i) A functional bean's methods do not by default allow for data persistence. An application or a business-logic programmer may choose to provide for explicit data persistence by programming the appropriate code, for example, by providing appropriate 15 JDBC statements.

Referring to the drawings, Figs. 1(a)-(f) depict several exemplary architectural views of a client transaction utilizing a functional bean. A client 110 executes a business method resident on an enterprise application server 100 via a network 102. It is assumed that the enterprise application server 100 comprises a container 120, which provides access to the 20 several kinds of beans—session beans 150, entity beans 160 and functional beans 180—of possibly several types, each type implementing at least one function—to access with a data store such as the database 130. External entities such as clients or other beans can access the EJBs via instances of EJBObjects. Fig. 1(a) shows the client 110 accessing session bean 150, entity bean 160 and functional bean 180 via EJBObjects 151, 161 and 181 respectively.

25 It is further assumed that there are different types of each kind of bean. For sake of simplicity, the details of how “handles” to the various types of beans are obtained by the client is not discussed here.

In a first aspect, shown in Fig. 1(b), the client 110 accesses a functional bean 180 and 30 executes the business logic via the EJBObject 181. The functional bean 180 in turn provides access to the database 130 as shown. In this scenario, a three-level interaction takes place: Client --> functional bean --> database view(s) or record(s) or any limited resource. The

container 120 includes the code to manage a client connection.

In a second aspect, shown in Fig. 1(c), a client 110 may invoke a session bean 150 via EJBObject 151 for connection management. The session bean 150 in turn invokes a functional bean 180 via EJBObject 182 to execute business logic and perform any database transactions as shown.

Referring to Fig. 1(d), a functional bean 180 manages the connection as in Fig. 1(b), but instead of providing access to the database 130 as in Fig. 1(b), invokes an entity bean 160—via an EJBObject 162—with a primary key or other identifier to interface with the database 130.

Referring to Fig. 1(e), another scenario includes a client 110 accessing a session bean 150 via an EJBObject 151 to establish and manage a session with the client 110. The session bean 150 invokes or mediates access to a functional bean 180 (via an EJBObject 183), wherein the functional bean 180 implements the business logic functions. The functional bean 180 invokes or utilizes the services of an entity bean 160 (via an EJBObject 163) to perform functions such as database updates. In a subsequent session, a new session bean 150 may be present to invoke the functional bean 180, but transactional persistence is taken care of by the container 120.

Referring to Fig. 1(f), a functional bean 180-1 of one type (type A) advantageously invokes a functional bean 180-2 of a different type (type B) (via an EJBObject 184) to accomplish a database transaction or any other type of transaction involving a limited resource.

Though the above describes a few scenarios, persons of ordinary skill in the art readily understand and extend the concept to arrive at other combinations of beans and bean types, all of which should be construed to be within the spirit and scope of the present invention.

#### IMPLEMENTING A FUNCTIONAL BEAN

In the following discussion, an implementer is defined as a business logic developer who implements the functional bean software for use by an application programmer.

**A FIRST EMBODIMENT: CREATING A FUNCTIONAL BEAN COMPILER AND CONTAINER**

Referring to Fig. 2, in a first embodiment, the functional bean 210 is derived from the EnterpriseBean class 204. The EnterpriseBean class 204 is advantageously derived from class java.io.Serializable 202. Other possible derivatives of the EnterpriseBean 2024 class include session bean 206 and entity bean 208. Referring to Fig. 1(a), in this embodiment, the functional bean 180 operates within a container 120, which provides for transparency to an application or business logic programmer.

The container 120 provides deployment tools that facilitate a deployer to supply "start-up" information for any external resources. The container 120 further involves the external resources in functional bean transactions if the resource is transactional and can support transactions. Additionally, the container 120 provides for resource pooling and manages the resource pool in a manner transparent to the functional beans 180.

A functional bean compiler can be used to compile such beans created within the application server. Persons of ordinary skill in the art know design and development of the functional bean compiler and container. This accords a native support for functional beans. A container with native functional bean support would automatically provide transparency to an applications programmer, thereby eliminating the need to deal explicitly with a Service Manager (not shown in Fig. 2). Further, developing a functional bean compiler and container will eliminate the need for using an entity bean to implement a functional bean.

Thus, in a first embodiment, a functional bean is implemented by enhancing an enterprise application server's container to support functional beans. In this case, an implementer (of functional beans) can incorporate the functions of a Service Manager (not shown) in the container 120, thereby making the Server Request Manager invisible to the application programmer.

**A SECOND EMBODIMENT: USING AN ENTITY BEAN TO SUPPORT A FUNCTIONAL BEAN**

In a second embodiment, an implementer may advantageously use an entity bean to develop the functional bean of the present invention. This embodiment enables users of commercially available EJB-compliant application servers who do not have access to EJB server's source code to modify it to support the new bean type. Such a user pays the price by losing the transparency of the new bean type. As an illustration, the Service Manager can be

advantageously created using the tools provided by a commercial EJB compliant container.

The Service Manager manages functionality such as functional bean pool management, configuration, and fail-over. Commercially available EJB-compliant container may be used to handle transaction management. The client, in collaboration with the Service Manager,

5 handles transactional persistence by requiring additional explicit method calls.

Thus, in the second embodiment, an entity bean is used to provide the functionality of a functional bean. In collaboration with the Service Manager, the client handles functional beans in a non-transparent manner. In this implementation, the implementer cannot bypass the default behaviors of entity beans such as data persistence because he cannot override the  
10 default behaviors associated with the methods behind the Application Programming Interfaces (API) of a third-party EJB server code. In this case, functional beans can be implemented as entity beans that contain no data attributes or only a few attributes. This selection helps minimize any overhead associated with persisting those data attributes.

Both in the first and in the second embodiments, the implementer or the application  
15 programmer can limit the total number of functional beans of each type to a limited number. This limitation may be based on a number of factors such as the availability of system resources. Thus, in the second embodiment, the database persistence of the functional beans may be a feature of using entity beans, but this feature could be used or not used advantageously based on a designer's suitable selections.

20

#### DEVELOPING AN APPLICATION USING A FUNCTIONAL BEAN

A functional bean provides a functional unit that utilizes one or more scarce resources—such as socket connections to a back-end system, or a memory intensive resource—for its functionality and hence needs to be shared across the clients that need such  
25 resources. The number of functional bean instances can be determined by the number of scarce resources available to the application and can be configured by an application deployer. Load balancing multiple client requests across available functional bean instances can be performed by the container or Service Manager.

A container manages the functional bean's life cycle. The container notifies the bean  
30 to perform some task and additionally achieves advantages such as scalability. Transactions can either be container managed or bean managed.

In an embodiment, the Java™ programming language or other similar method may be used to develop an application using the functional bean. In both the first and the second embodiments, the invented functional beans are developed by (1) writing the code for home interface; (2) writing the code for remote interface; (3) writing the code for functional bean class; (4) compiling the code; (5) creating a deployment descriptor, which contains information about the functional bean including environment such as variables, references to other beans, and resource factories; and (6) creating a standard packaging jar file format, the “ejb-jar” file.

As an example, the functional bean may have the following remote interface.

```
10      import java.rmi.RemoteException;
11      import javax.ejb.*;
12      public interface OrderEntry extends EJBObject {
13
14          // list business methods here.
15          public String count() throws RemoteException;
16      }
17
18      Illustratively, the Home Interface can be described as follows.
19
20      import java.rmi.RemoteException;
21      import javax.ejb.*;
22
23      public interface OrderEntryHome extends EJBHome {
24          public OrderEntry create (String customerName, String customerAccount)
25              throws RemoteException, CreateException, BadAccountException;
26          // Remove and other create methods go here.
27      }
```

### 30 RESOURCE POOLING

A functional bean can access resource factories, i.e., it can access databases, or pools of connections associated with any other type of resource. A deployer identifies the resources a functional bean needs and provides references to the resource factories in the

deployment descriptor. In general, the deployment descriptor declaration includes description, name, type and any authorization required to access the resource factory.

Referring to Fig. 3, the method by which a functional bean 180 obtains a handle to an external resource is depicted. In general, a provider of functional beans specifies the attributes of any external resources a bean needs during its lifetime. The functional bean 180 looks up a resource name using JNDI 310 to obtain an instance of an external resource such as a data source 130. This instance of the external resource interface 320 is used to access resources, which can be used by a functional bean 180 subsequently. If the functional bean 180 is configured to support bean-managed transactions, the functional bean 180 invokes a “begin” method 345 in the external resource interface 320 to transmit a message to the container 120 indicating that the functional bean 180 intends to begin a transaction that involves the external resource interface 320. After the begin method executes, the functional bean 180 performs a method 347 such as UpdateOrderEntry. After control of execution returns from the method 347 call, the functional bean 180 may release any external resources that it has acquired by calling a “close” 355 method on the external resource interface 320 after appropriately committing/rolling back the transaction. In an embodiment, in container-managed transactions, the functional bean does not need to explicitly begin and commit or rollback a transaction. This could be handled by the container 120 in a manner transparent to functional bean 180 before and after invoking a method on the functional bean 180. The following piece of code illustratively implements the method shown in Fig. 3.

```
public class LegacyGatewayBean implements SessionBean {  
    EJBContext ejbContext;  
    public void talkToExternalApplication(...) {  
        ...  
        // Obtain the initial JNDI context  
        Context initCtx = new InitialContext();  
  
        // JNDI lookup for the external resource factory  
        com.foo.bar.ExternalDataSource eds =  
            (com.foo.bar.ExternalDataSource) initCtx.lookup(  
                "java:comp/env/external/LegacyConnectionPool");
```

//Invoke the factory to obtain a resource  
com.foo.bar.ExternalResource res = eds.getResource();  
  
...  
}  
5     }  
  
// The external resource interface could be a simple interface in order to  
// accommodate diverse external applications. It defines the methods that signal  
// the container of the bean's intended transactional boundaries. Implementers  
10    // of external resources can extend the ExternalResource Interface to provide  
      // methods that address external application-specific functionality.

public interface ExternalResource {  
    void begin();  
    void commit();  
    void rollback();  
    void close();  
}

20    A functional bean deployer identifies the resources a functional bean may require  
during its lifetime and supplies mapping information needed to start-up to the external  
resources. The deployer binds the resource factory reference to a resource factory that exists  
in the target environment such as a database. A deployment descriptor is a device that allows  
the deployer to make this binding.

25

The following code sample illustrates how the deployer and the deployment  
descriptor achieve their respective functions.

30    // Specifying resource references in a functional bean's deployment descriptor.  
<enterprise-beans>  
      <session>  
          <ejb-name>LegacyGateway</ejb-name>  
          <ejb-class>com.foo.bar.LegacyGatewayBean</ejb-class>  
      ...

D E C I D E D  
D E C I D E D  
D E C I D E D  
D E C I D E D  
D E C I D E D  
D E C I D E D  
D E C I D E D  
D E C I D E D

```
<resource-ref>
    <description>
        // data source for database that stores information needed by Legacy Gateway bean.
    </description>
    5          <resource-ref-name>jdbc/LegacyGatewayDB<res-ref-name>
    <res-type>javax.sql.DataSource</res-type>
    <res-auth>Container</res-auth>
    </resource-ref>
    <resource-ref>
    10         <description>
        // the data source for the pool of connections to the Legacy application.

    ...
    </description>
    <res-ref-name>external/LegacyConnectionPool<res-ref-name>
    15          <res-type>com.foo.bar.ExternalDataSource</res-type>
    <res-auth>Container</res-auth>
    </resource-ref>
    </session>
    </enterprise-beans>
    20
```

#### AN EXAMPLE APPLICATION: A TROUBLE ENTRY SYSTEM SOLUTION

An illustration of the disclosed principles is described in detail in the following with reference to an exemplary system for entering trouble tickets in a telephone company, which system is named “Trouble Entry System Solution” (TESS). It should be noted that this is used only to illustrate the principles of the present invention and not as a limitation.

Referring to Figs. 4(a)-(b), a hardware architecture may comprise a plurality clients 110 coupled to a network 102, to which is communicatively coupled an enterprise application server 100 incorporating the inventive principles disclosed herein. A database 130 is communicatively coupled to the enterprise application server 100. In an embodiment, the database 130 is a relational database management system such as the Sybase RDBMS. In other embodiments, the database 130 could comprise a number of legacy database systems running on different servers (not shown) communicatively coupled to the enterprise application server 100 via data communication protocols such as JDBC, TCP/IP sockets or others (not shown).

In an embodiment, the enterprise application server 100 comprises a programmed general-purpose computer including one or more central processor unit(s) (CPU) such as a Sparc™ processor marketed by Sun Microsystems, Inc. of Palo Alto, California. Further, the enterprise application server includes a storage device such as a hard disk, a memory device such as a semiconductor memory, a communications device such as a network card to connect the enterprise application server 100 to the network 102. The enterprise application server additionally comprises an operating system such as Solaris™ or Linux and is configured to communicate with other computers coupled to the network 102 via a protocol such as the Transmission Control Protocol/Internet Protocol (TCP/IP).

10 In addition to the configuration described above, the enterprise application server 100 may include server software based on the Enterprise JavaBeans (EJB) standard. Illustratively, the invented functional beans 180 are implemented by modifying the entity bean objects provided in the WebLogic™ (formerly, Tengah™) marketed by BEA Software Corporation of San Jose, California.

15 In alternative embodiments, the enterprise application server 100 is designed to support both Java application-based clients as well as web browser-based clients. The basic difference between the two is the location of the user presentation logic, and the type of components used to implement them. The Java application clients contain the presentation logic layer, while the web-based clients contain no presentation logic beyond the functionality provided by the web browser: the whole presentation logic layer is contained in 20 servlets residing in the browser. The first can be described as "fat" clients, while the latter are known as "thin" clients.

25 Referring to Fig. 5, which shows an exemplary architecture of software 500 executing on the enterprise application server 100. The software architecture 500 comprises a plurality of layers of software, each layer including a separate subsystem, for example, a view layer 502 (which includes the presentation logic); an application-model layer 504 (which is the middleware layer incorporating the business functions); a domain layer 506 (comprising business data persistence); and a persistence layer 508 (which includes code to accomplish data persistence). The layers are shown horizontally and the subsystems 522, 524 and 526 30 are shown vertically. Each layer includes a set of classes such as Enterprise JavaBeans that share certain responsibilities, utilize the services provided by a layer below, and/or provide

services to a layer above.

The subsystems run across all layers. All the classes including the beans in a subsystem implement one or more related use cases, and thus provide a clearly defined and cohesive subset of the functionality. Persons skilled in the art understand how to construct  
5 use cases from a user requirements definition, and therefore, that method is not elaborated here. Each class corresponds to an individual, clearly defined piece of functionality. Examples of such functions include trouble entry and circuit testing. Subsystems are designed to be independent of each other and each subsystem is configured to satisfy specific business requirements.

10 In addition to the APIs provided, each layer is configured to include a set of common functions—known as frameworks—shared by all subsystems in that layer. In the example, these are configured to provide functions such as user authorization and authentication, metrics, and database access service with transparent fail-over, and the like.

15 As an illustration, in one embodiment, a DatabaseService component is configured to allow client objects from several layers to access the database without regard to the details of connection setup and management. In alternative embodiments, the database 130 could include a first database (a primary database), and a second database (a backup database), with a method that enables fail-over from the primary database to the backup database. The client objects are not aware from which pool they get their connections, or even which database is  
20 designated as the primary database.

Further, instances of classes—whether within the same layer or those from different layers—communicate via messages. This method of communication via messages is called coupling. An embodiment achieves weak coupling across layers by minimizing the number of classes from each layer that communicate with classes from other layers. In this  
25 architecture, database access, connection pooling, and database fail-over subsystems are weakly coupled to all other components.

Referring to Fig. 6, an EJBObject 602 is configured as a network-visible, remote object, which prevents direct access to an instantiated functional bean 180 and acts as a stand-in for the functional bean 180 that has restricted access. Advantageously, the  
30 functional EJBObject 602 has a unique identifier that is assigned by the container 120. Further, the container 120 maintains a handle to the EJBObject 602, which can be used to

access the functional bean 180.

Referring again to Fig. 5, when EJBs-functional, entity, and session beans-are used to implement subsystems in a given layer they are advantageously organized to provide a weak coupling between that layer and the one immediately above it because

5 they provide a facade. Objects in the upper layer do not send messages directly to a non-adjacent lower layer's beans but to their interface objects, i.e., their respective EJBObjects (not shown in Fig. 5) through the intermediate layers. These objects act as proxies to the true EJBs, providing a simplified, controlled interface to the external world, isolating those EJBs from direct exchanges with external objects.

10 In the enterprise application server 100, several types of beans implement the business logic. In most cases, a single functional bean implements one specific business function, such as trouble ticket entry, history, transfer, and status. In some embodiments, other types of functional beans are also contemplated to work together with one type of functional bean. For example, the Trouble Entry functional bean also invokes the Trouble Status functional as part of its function. On the other hand, the Trouble Transfer functional bean does not call any other functional bean or EJB. The choice of using or not using other functional beans depends on a particular function to be implemented and the amount of code redundancy introduced by strictly adhering to developing functional beans that are independent of each other, i.e., loosely coupled. Functional beans that are independent of  
15 each other are attractive because they incorporate the virtues of loose coupling between subsystems; they make it easy for programmers to work independently and in parallel. This simplifies design, coding, unit and integration testing.

20

Some functional beans, such as TroubleEntry and TroubleTesting, are implemented as entity beans because they use scarce system resources such as the total number of socket connections to an external system such as a database. A system administrator or other deployer of beans determines the maximum number of functional beans of each type. The system administrator can set this maximum number not to exceed the total number of limited resources such as the number of simultaneous database connections available.

Functional beans are used to model functional (not data) entities, such as tellers in a bank. As customers arrive at a bank, i.e., as client requests arrive at the enterprise application server 100, a Service Manager routes the requests to an available teller from a

pool of tellers. Instances of functional beans model the particular type of transaction the client requests-trouble ticket entry, account history, account status, and others. Thus, the Service Manager (the bank's concierge) controls access to the functional beans as described above, thereby achieving the following interaction:

5 Client → Functional bean → Database view/record

Referring to Fig. 7, the steps included in a typical transaction are shown. As a first step, functional beans are created to model business service providers (step 700). A special functional bean called a Service Manager is also created. Then, suppose a client 110 requires a service provided by a particular type of functional bean, say, Trouble Ticket Entry. The 10 client 110 first calls the Service Manager (step 702). Responsive to this request, the Service Manager initializes a handle to an instance of the Trouble Ticket Entry functional bean (step 704). The Service Manager sends the handle to the client 110 (step 706).

Thereafter, the client 110 can directly contact the Trouble Ticket Entry functional bean in the following way. The client 110 issues a request-for example, a request-for a 15 transaction (step 708). This request is queued by the Trouble Ticket Entry bean (step 710). By queuing the requests, the clients do not get a response such as an "I-am-busy" exception from the TroubleTicketEntry bean whenever it is busy. Later the Trouble Ticket Entry bean executes the client's transaction (step 712).

At any time the client is free to access any number of functional beans of a particular 20 type or of many different types. Adding additional clients causes the Service Manager to invoke a load-sharing program to instantiate additional functional beans based on the number of available resources. Actual clients will be assigned to existing functional EJBs using a load-balancing algorithm such as round robin. This makes the functional bean architecture fully scalable.

25 Since a functional bean can be developed as a transactional object, it can participate in a distributed transaction, i.e., calls to transaction-aware beans can be chained with full two-phase commit and rollbacks spanning several functional beans.

It should be noted that if a particular service does not need a functional bean's unique 30 features, a developer may choose to implement the service as a stateful bean, a stateless bean, an entity bean, a java object, a java class, or even an extra method on an already existing java class. In alternative embodiments this choice is made based on the level of functionality,

transactional-awareness and persistence required for that service. When services are implemented as session beans, it usually follows the following three-tiered model:

Client → session.bean → External service

Or,

5 Client → session bean → Database view/record

Referring now to Fig. 8, in an alternative embodiment, a web-based client 110', which is a client such as the client 110 described above, except that it provides a browser program as a method to access the services offered by the enterprise application server 100 via a network 102 such as the Internet. An example of a browser program is the Internet...

10 Explorer marketed by Microsoft Corporation of Redmond, Washington. In this embodiment, the enterprise application server 100 is additionally configured to function as a web server. A web server typically serves web pages to the client, which web pages allow a user to send requests for information and receive response messages in the form of packets that display the information in the browser web pages. Additionally, in this embodiment, servlets 810 are  
15 advantageously used to provide the presentation logic for the user. Servlets 810 are server side components that enable a part of the presentation logic to be executed on a web server. They also can be configured to provide an interface to back-end components 820, or an external service such as a database 130. These back-end components 820 could comprise EJBs such as entity, session and functional beans that provide business services directly to  
20 application-based clients as described above. Thus, as shown, both the web-based client 110' and application-based clients 110 share functional beans as back-end components 820. This embodiment can also be implemented using Java Server Pages (JSP), which provide functionality equivalent to servlets. Server-side Common Gateway Interface (CGI) scripts and client-side applets or JavaScript can also provide alternative implementations, amongst  
25 other ways that are functionally equivalent to servlets.

The foregoing load-balanced and scalable access to limited system resources describes a new and useful enterprise software component that allows transactionally transparent persistence to clients over a communication network. Persons skilled in the art may make several modifications, enhancements or rearrangements without departing from  
30 the spirit and scope of the invention or without undue experimentation. For example, the invented software component can be implemented in a CORBA or a DCOM-compliant...

server; the enterprise application server 100 described herein can be scaled using a method such as clustering; several enterprise application servers may be configured to specialize a single or multiple functions to achieve a distributed functional architecture; resource management can be distributed; garbage collection can be provided for; high availability can  
5 be achieved by using standard methods; the object oriented components can be made reusable; complex objects can be transparently exchanged between components with a need to create object interface descriptions such as the IDL files required in CORBA; and clients can be implemented using any one of the known methods such as applet-servlet, or applet-RMI-EJB. Accordingly, all such modifications, enhancements, rearrangements and  
10 departures should be construed to be within the spirit and scope of the appended claims.

CONFIDENTIAL